

Implementing the HP adaptive infrastructure using ProLiant technologies

technology brief



Abstract.....	2
Introduction.....	2
Vision for an adaptive infrastructure.....	3
Core technologies within ProLiant platforms	4
Continuous and secure operations.....	4
Server and component fault management.....	5
Performance issue identification	5
ProLiant failover clustering technologies	6
Dynamic internet solutions architecture.....	6
Automated, intelligent management.....	6
Lights-Out management technologies	8
Other system-level management tools.....	9
HP OpenView.....	9
Dynamic resource optimization.....	10
Scale-up architectures	10
Scale-out architectures.....	11
Workload optimization	12
Virtualization	12
Resource deployment and provisioning	13
Open industry standards	14
Hardware architectures.....	14
Fabric interconnects	14
Software interfaces.....	15
Integrated with solutions and services.....	15
Adaptive infrastructure services.....	15
Adaptive infrastructure solutions.....	15
Collaborative industry partnerships	16
Conclusion.....	16
Call to action	17

Abstract

Businesses of all sizes are continuously challenged to compete in today's fast changing, global economy. Resources are scarce and business models are constantly transforming, while customer expectations continue to increase. To satisfy these business demands, information technology (IT) professionals must create an infrastructure that adapts easily to change, conserves resources, and is responsive to customer requirements. This paper describes how the HP vision for an adaptive infrastructure can be implemented using ProLiant servers and industry-standard technologies to meet those IT challenges. It describes three interrelated technology areas that work together to form the basis of an adaptive infrastructure: continuous, secure operations; automated, intelligent management; and dynamic resource optimization. HP's strategy for delivering the adaptive infrastructure builds on a rich heritage of developing innovative products using open industry standards. The adaptive infrastructure is integrated with value-added solutions and services from HP and third-party providers that will enable businesses of all sizes to scale operations to meet changing IT demands. The paper gives examples of the value of ProLiant technologies in defining the adaptive infrastructure.

This paper assumes that the reader has some familiarity with existing ProLiant products and technologies.

Introduction

In today's always-on, global environment, it seems the only constant in the business world is the need for change. Business models are dynamic, with organizational structures being affected by mergers and acquisitions, uncertain economic environments, and markets that may expand or decline suddenly. Businesses must find ways to adapt quickly and improve efficiency or be left behind.

Because the environment is so challenging, companies must reduce IT costs and complexity while ensuring that enough resources, and the right resources, are allocated to business priorities. To thrive, a business must clearly prioritize its needs and find ways to conserve its resources.

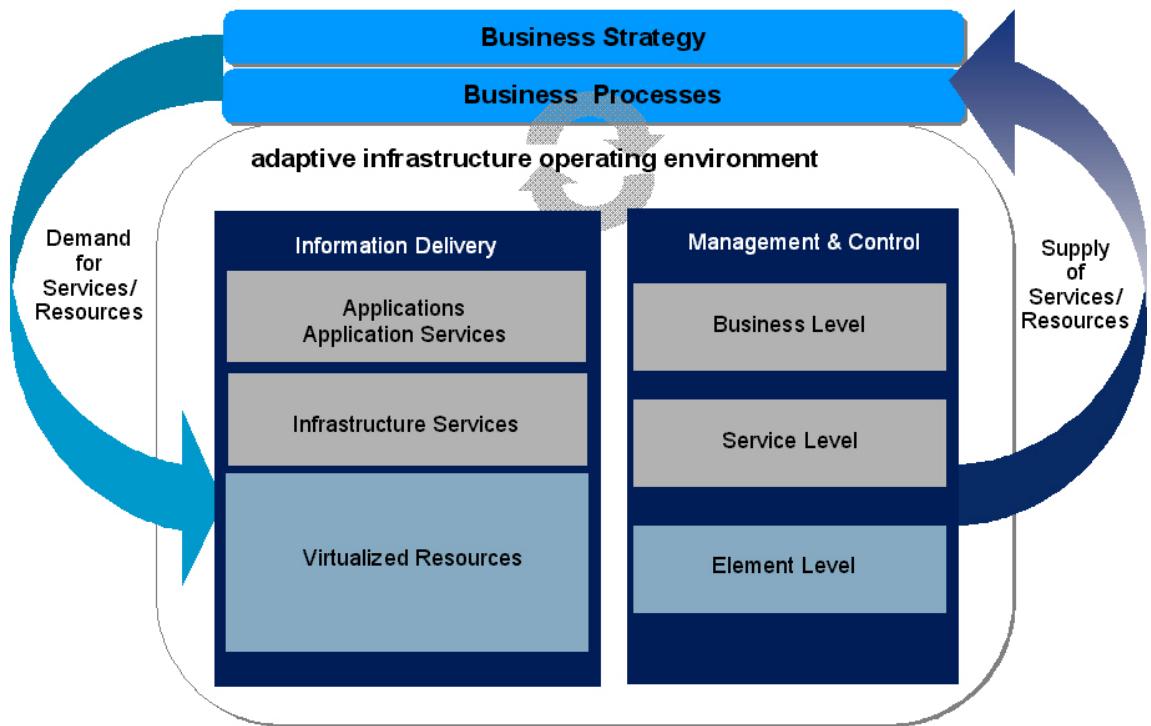
At the same time, customer expectations for service are high. Accustomed to the promise of a web-enabled world, today's customers expect business information and services to be constantly available, from anywhere, in the way they deem most suitable, and with a guaranteed quality of service. Therefore, IT organizations must virtually eliminate downtime and be able to respond immediately to customers. In many cases, such as in the service provider industry, customers are demanding service level agreements that compensate them for poor service or downtime.

The result of these challenges is a need for what HP calls "business agility." Business agility is the ability to respond nimbly to opportunities in the marketplace as well as respond effectively to change. It refers not only to how quickly a company institutes change, but also to how efficiently a company can change — how well risk is minimized and how many resources are spent implementing change.

Business agility requires an IT infrastructure that adapts as business and IT conditions change — an adaptive infrastructure. The technology infrastructure must become a core asset that enables, rather than inhibits, business growth.

HP's overall approach for the agile business starts with business strategy and processes, and using the Adaptive Infrastructure Operating Environment (AIOE), manages the IT infrastructure to enable and support those processes (Figure 1). AIOE depends on enterprise applications, infrastructure services, virtualized resources, and management and control throughout all levels: business, service, and system-level. As the AIOE is implemented, it enables IT staff to conserve resources, align those resources to evolving business needs, and anticipate and respond immediately to customer demands.

Figure 1. The adaptive infrastructure operating environment results in the ability to take immediate action based on business strategy and processes.



Vision for an adaptive infrastructure

The HP adaptive infrastructure tightens the link between IT and business objectives so customers can anticipate and rapidly respond to new market requirements. HP is developing the tools, building the platforms, architecting technologies, and initiating partnerships with other industry leaders to create a truly adaptive infrastructure. Most importantly, HP's broad portfolio of technologies, solutions, and services is aligned to the adaptive infrastructure vision of an efficiently linked IT and business infrastructure. Individual HP products and integrated solutions for the adaptive infrastructure are ready for implementation today.

The adaptive infrastructure is powered by three interrelated, core capabilities that are woven throughout HP server, software, and storage products:

- Continuous and secure operations — the ability to establish and maintain required levels of availability and security throughout the infrastructure. Maintaining such operations includes the ability to predict, diagnose, and respond rapidly to potential and actual fault conditions. This capability provides a highly stable infrastructure, beginning with individual server subsystems and extending across the enterprise, ensuring business continuity while minimizing the cost of downtime.
- Automated and intelligent management — the ability to plan, monitor, control, and initiate responses to changing demands, based on established service level agreements or rules. This management ability includes the control of hardware resources as well as software deployment, maintenance, and provisioning tasks. Automated and intelligent management reduces complexity in the datacenter by allowing IT staff to focus on strategic business planning rather than on routine administration tasks.

- Dynamic resource optimization — the ability to provision, assign, and share technology resources when needed, wherever needed, based on business or operational requirements. This capability provides greater flexibility in the utilization of resources.

HP's vision of an adaptive infrastructure includes hardware and software based on industry-standard technologies, strengthened by collaborative industry partnerships, and highly integrated with solutions and services.

The result will be an adaptive infrastructure that enables:

- Business decisions to be informed by a more comprehensive and immediate view of all core business operations.
- Optimal use of computing, network, software, and storage resources, enabling people to be more productive and use assets more effectively.
- Increased operational efficiencies and uptime, leading to cost savings and greater customer satisfaction.

Core technologies within ProLiant platforms

The ProLiant platform focuses on delivering innovative technology based on unifying industry standards. This strategy has successfully driven down the economics of enterprise computing while pushing standards-based servers higher into the enterprise. HP is applying the same approach to deliver HP's adaptive infrastructure on ProLiant platforms.

When using ProLiant platforms that integrate the three core capabilities, a customer can attain the full power, responsiveness, and efficiency of an adaptive infrastructure, while maintaining a low cost of ownership. For example, using intelligent management technologies and automated provisioning tools, an IT administrator can rapidly deploy servers globally with a minimal amount of effort. Once they are configured, integrated hardware and software can minimize system downtime through redundancy and error-correcting mechanisms. Software updates can be automatically downloaded according to the customer requirements. Administrators can keep resources continuously available by using management and monitoring tools that automatically identify and resolve fault conditions. When conditions change, the administrator can pinpoint areas that need to be scaled up or down in response to changing business needs.

Developing an adaptive infrastructure is an incremental process. Using ProLiant technologies, customers can build out the adaptive infrastructure using the products, solutions, and comprehensive service offerings that fit their requirements. Then they can scale up and scale out as needed to broaden the scope of their adaptive infrastructure.

The following sections give representative examples of how HP products fulfill the vision of each core technology in the adaptive infrastructure. Some products or solutions cited here make use of more than one core technology. For example, the next-generation ProLiant servers include the ability to scale dynamically and provide intelligent fault management.

Continuous and secure operations

HP offers a broad range of business continuity and high-availability solutions to address a customer's need for stable, continuous operations across the entire infrastructure. At the system level, ProLiant servers have built-in redundancy, hot-plug components, and fault-tolerant technologies. Using tools such as Insight Manager 7, an administrator can perform routine performance checks on ProLiant servers to identify and understand server bottleneck issues. ProLiant servers can be cost-effectively grouped into clusters and combined with a range of HP shared-storage products to provide capabilities ranging from highly available clusters to massively parallel clusters. Moving beyond

clusters to enterprise-wide architectures, customers can develop a highly available Internet environment using architectures such as the HP dynamic internet solutions architecture (DISA).

Other core capabilities of the adaptive infrastructure — automated intelligent management and dynamic resource optimization — increase an administrator's ability to identify problems before they happen, to identify other resources that can be used, and to dynamically shift resources around potential problems. The ultimate result will be an infrastructure that provides continuous operations across the entire enterprise.

Server and component fault management

To reduce downtime, HP has introduced many features that improve availability in server and storage subsystems. HP has delivered systems for many years with innovative hot-plug components such as drives, fans, power supplies, and PCI slots. Each hot-plug component reduces unplanned downtime by allowing a customer to keep the system up and running while replacing a failed component. For example, when using Smart Array controllers with Advanced Data Guarding (RAID¹ ADG), a system can tolerate two simultaneous drive failures without downtime or data loss. Because this advanced RAID level allows a system to sustain a second drive failure while rebuilding a failed drive, RAID ADG offers dramatically higher fault tolerance than RAID 5. This capability is extremely valuable when compared to the time required to rebuild today's ever larger drive capacities.

HP has added another dimension to server fault management with the introduction of Hot Plug RAID memory.² Hot Plug RAID Memory provides key functionality that has previously been unavailable within industry-standard servers. Like redundant array of independent disk technology used in storage subsystems, Hot Plug RAID memory uses a redundant array of industry-standard DIMMs³ to provide fault tolerance and the ability to hot replace and hot add memory while the server is operating.⁴ With Hot Plug RAID Memory, the RAID engine provides redundancy to ensure data protection, and the hot plug abilities allow administrators to replace a DIMM without any downtime.

Within the modular ProLiant BL server blade line, intelligent components work together to ensure that fault alerts are forwarded to the appropriate modular component. For example, the ProLiant BL p-Class system⁵ isolates the central processing units (the blades) from the power subsystem to facilitate rapid deployment. If any power issues occur, the power subsystem and several other intelligent components forward appropriate alerts only to the affected blades, ensuring availability of the overall system.

Performance issue identification

Performance tuning and optimization is inherent to good infrastructure management. While availability technologies such as fault-tolerance are reactive in nature, the ProLiant Essentials Performance Management Pack (PMP) is a proactive tool that helps an administrator identify bottlenecks on a server before they become critical issues. PMP examines the values of certain hardware and operating system (OS) performance counters and applies a series of algorithms to analyze the data. Then the administrator can identify if a performance bottleneck exists in any of the servers' subsystems: processor, memory, network I/O, disk I/O, or PCI bus. Performance issues are analyzed and can be debugged down to the individual component level, such as a single disk drive.

¹ RAID is an acronym for Redundant Array of Independent Disks, a common method of fault tolerance using multiple disk drives configured into a single logical array.

² For more information on Hot Plug RAID Memory, see the technology brief titled "Hot Plug RAID Memory technology for fault tolerance and scalability," available at <http://h200001.www2.hp.com/bc/docs/support/SupportManual/c00257001/c00257001.pdf>.

³ DIMM is an acronym for dual inline memory module.

⁴ Hot replace capability will be available out of the box, with no need for any specific operating system or driver support. Hot add and hot upgrade capabilities require support from the operating system to recognize the additional memory that is available. HP is working with the major operating system vendors to ensure these capabilities will be supported in upcoming releases.

⁵ For more information about the ProLiant BL p-Class system, see the website at www.hp.com/products/servers/proliant-bl/p-class/info.

PMP is integrated with Insight Manager 7 so that administrators can assign events and configure notification tasks for performance events just as they would for hardware component failures. If PMP warns of a heavy workload in one of the monitored subsystems, the administrator can review the logged data file and adjust workloads or configurations. Thus, the system administrator can routinely monitor and adjust workloads to achieve higher throughputs and lower response times.

ProLiant failover clustering technologies

As a customer requires more availability, HP provides a wide range of industry-standard clustering solutions.⁶ Customers can choose from the entire scope of ProLiant servers, allowing them to match the processing requirements of their applications to the level of availability and performance needed.

For entry-level or remote site cluster requirements, the ProLiant DL380 Packaged Cluster, consisting of two ProLiant server nodes and Smart Array Cluster Storage⁷, greatly simplifies implementation. The servers include high-availability features such as Online Spare Memory, redundant power supplies, fans, and shared-storage controllers. Another solution for entry-level availability is the ProLiant Essentials Recovery Server Option Pack. This software allows an active-standby configuration of two servers in which one server acts as the primary or active server, and the second is the recovery server in a passive mode. When a fatal fault is detected in either the primary server or its OS, the recovery unit will initialize and take over.

For larger or more complex clustering requirements, ProLiant Clusters support clustering software for Windows, Linux, Oracle 9i RAC, and Novell. These cluster solutions provide application or data failover as well as improved scalability and performance. ProLiant servers can be combined with flexible storage options such as the SCSI-based Smart Array Cluster Storage, business-class fibre channel MSA1000, or enterprise level Enterprise Virtual Array, to provide storage scalability from gigabytes to multiple terabytes.

Dynamic internet solutions architecture

The DISA architecture⁸ provides the framework for building an adaptive infrastructure in web-based environments. DISA defines three primary computing layers: the access and acceleration resources layer, the web and application resources layer, and the data resources layer. ProLiant servers are ideal for this multi-tier environment, because they are designed with highly available subsystems and fault tolerance built in. In addition to the inherent advantages of the ProLiant servers, the DISA structure includes redundancy to minimize or eliminate potential single points of failure. The DISA structure, along with the best practices that HP provides, results in an extremely stable infrastructure with reduced risk of unplanned downtime for a wide range of e-commerce environments.

Automated, intelligent management

Automated, intelligent management, the second core capability, encompasses the ability to monitor, maintain, and automate the control of hardware resources. It also includes the ability to deploy, provision, and maintain software automatically and intelligently. HP's broad range of system-level to service-level management capabilities enables comprehensive service, performance, and fault management (Figure 2).

⁶ More information about ProLiant clusters is available at: www.hp.com/servers/proliant/highavailability.

⁷ For more information on Smart Array technology, see www.hp.com/products/smartarray.

⁸ For more information, see <http://activeanswers.compaq.com/ActiveAnswers/Render/1,1027,1573-6-100-225-1,00.htm>.

Figure 2. Examples of system-level and service-level management tools by HP



Automated, intelligent management tools take the cost and headache out of routine maintenance, providing the IT administrator with superior capabilities to deliver near-continuous uptime and ensure target IT service levels.

At the system level, the ProLiant Essentials Foundation Pack Software (which includes Insight Manager 7, Management Agents, and SmartStart) provides the basis for automated, intelligent management. Optional tools such as Integrated Lights-Out, Remote Insight Lights-Out Edition II, the ProLiant Essentials Rapid Deployment Pack, and the ProLiant Essentials Workload Management Pack build on that foundation to enhance system-level management capabilities. At the service level, heterogeneous management tools such as HP OpenView enable an administrator to monitor and control resource health, track utilization, and report on how infrastructure operations may impact the overall business. When Insight Manager 7 and OpenView are linked together, they provide the optimum closed-loop management tool from the system to the service level.

ProLiant Essentials Foundation Pack

Every ProLiant server ships standard with the ProLiant Essentials Foundation Pack. The software provides the tools needed to set up, configure, and manage ProLiant servers, including:

- Management Agents that monitor a server's health through the extensive hardware instrumentation built into ProLiant servers. Management Agents monitor parameters on system and subsystem elements such as disks, processors, memory, fans, and temperatures.
- SmartStart (included for ProLiant ML and DL servers) that simplifies setup and configuration of ProLiant servers.
- Insight Manager 7 that provides both reactive and proactive alerting to maximize system uptime and reduce the time required for infrastructure maintenance.

- ActiveUpdate that provides proactive notification and delivery of the most recent software updates, simplifying software maintenance.

Insight Manager 7

Insight Manager helps reduce the cost of managing critical IT assets and improves a system administrator's efficiency by allowing rapid responses to problems, maintaining consistent system software baselines across groups of ProLiant servers, and providing direct access to the remote administration capabilities delivered through Lights-Out management devices. Insight Manager 7, in combination with the Management Agents, monitors all server subsystems to ensure that they are running properly. Administrators can receive e-mail and pager notification of events and access to management information via a web browser, reducing the need for personnel dedicated to monitoring a fixed management console. For ProLiant blade servers, Insight Manager 7 provides a visual representation of the blade server environment to properly identify the location of each blade server within its enclosure and rack. Insight Manager 7 simplifies inventory reporting capabilities, dramatically reducing the time and effort needed to manage computing assets, from servers to laptops.

Future enhancements to Insight Manager will include the ability to manage from a Windows, Linux, and HP-UX workstation. Customers will be able to install Insight Manager on any of those platforms, providing a breadth of management capabilities that was previously unavailable within Insight Manager 7.

Active Update

Administrators are continually challenged with maintaining software updates to keep the infrastructure current and secure. Active Update provides proactive notification and delivery of HP software updates based on a customer-provided profile. After this notification, the administrator may review the updates to determine which updates should be applied and whether an individual server should be updated. Insight Manager 7 streamlines and automates the process by allowing the customer to define policies for automatically updating drivers and agents and to establish a baseline policy for each server. After the baseline is established, the administrator can request version control reports for a single server or multiple servers. Through Insight Manager 7, the administrator is able to update multiple servers automatically and remotely using a single software update task. By automating as much of the update process as possible, administrators can implement updates quickly, reduce human error, and reduce the potential for security breaches or other issues because of outdated software. The result is a more stable and secure infrastructure with less downtime.

Lights-Out management technologies

HP gives customers freedom from the server console through the Remote Insight Lights-Out Edition II board and the Integrated Lights-Out management processor (embedded in specific ProLiant servers). Since its introduction in 2000, ProLiant Lights-Out technology has provided customers the means of extending remote server management far beyond traditional emergency response. Lights-Out technology gives IT administrators the ability to control a Lights-Out enabled server from anywhere, at anytime, from multiple types of access devices, using a standard web browser – regardless of the condition of the OS or the server itself.⁹

Integrated Lights-Out (iLO) and Remote Insight Lights-Out Edition II are closely associated with Insight Manager 7, so that alerts can be forwarded to Insight Manager 7. Administrators can perform a query in Insight Manager 7 on all Lights-Out technology devices (known as management processors), allowing the administrator to have a single web page linking to all Lights-Out technology devices (see Figure 3). As an example of the benefit of Lights-Out technology, iLO is used in the ProLiant BL p-Class System to provide an intelligent communication channel to coordinate events, alerts, and location

⁹ For more information about Lights-Out technology, see the website at www.hp.com/servers/lights-out.

data throughout the modular server blade infrastructure (the server blades, the server blade enclosure, the power supplies, the power supply enclosure, and the power distribution). This intelligent communication ensures robust system management capabilities across multiple blades and multiple enclosures within a rack.

Figure 3. Example of how Insight Manager 7 can list all management processor devices

The screenshot shows the HP Insight Manager 7 SP2 interface. At the top, there's a toolbar with various icons. Below it is a header bar with the HP logo, the title 'Insight Manager 7 SP2', and sections for 'device status' (showing counts for 1, 3268, 1758, 139, 3315, 1787) and 'uncleared events'. A timestamp 'last update: Thursday January 9, 2003 - 10:46:10 AM' is also present. On the right side of the header is a 'support | logout' link. The main menu bar includes 'Home', 'Devices', 'Tools', and 'Settings', with 'Devices' being the active tab. To the left, a sidebar navigation tree includes 'Overview', 'Tasks', 'Reports', 'Queries' (which is expanded to show 'All Queries', 'Device' (with 'All', 'Devices by Type', 'Devices by Status', 'Devices by Operating System', 'System functions', 'Event', 'Cluster'), 'My Favorites' (with 'Configure', 'Singapore Data Center', 'Houston Data Center', 'San Paulo NDC'), and 'Logout'. The central content area is titled 'Query Results: Remote Insight Boards'. It features a table with columns: HW, MP, SW, PF, Device Name, Device Type, Device Address, and Product Name. Two rows are listed: one for 'aspen9r in Server im7tes' (Management Processor, 170.50.1.28, Remote Insight Lights-Out Edition II) and another for 'rib10m in Server 10m' (Management Processor, 170.50.1.51, Remote Insight Lights-Out Edition II). Below the table, a status bar displays 'Devices in table: 0 Critical, 0 Major, 0 Minor, 2 Normal, 0 Unknown. Total: 2'. A 'HELP' icon is located in the top right corner of the results window.

Other system-level management tools

HP has leveraged the Lights-Out technology to develop a management console for the ultra-dense ProLiant BL e-Class System. The ProLiant BL e-Class Integrated Administrator integrates with existing Insight Manager 7 and Simple Network Management Protocol (SNMP) tools while providing specific functionality for the ProLiant BL e-Class System. This functionality includes local and remote access to the server blades through a text-based console, the ability to power on and off the server blades remotely, and health monitoring of the server blades and enclosure.

For servers without Lights-Out technology, the HP Keyboard Video Mouse (KVM) IP Consoling Solution provides the ability to manage KVM switches over any internet protocol (IP) link. The intelligent KVM switch supports up to three network users, who can access applications and screens remotely. A java applet installed on a Windows or Linux workstation provides network users with high-performance KVM access and requires no additional software to be installed on the server. Pre-defined alert messaging can be sent to Insight Manager 7, HP OpenView, or other SNMP management software programs.

HP OpenView

Management tools like Insight Manager 7 identify fault and performance issues at the system level — before the issues become failures that would cause disruption and downtime. With integrated tools from HP OpenView, an administrator can manage heterogeneous datacenter environments containing

virtually any type equipment as well as services such as email, enterprise resource planning, and e-commerce. The administrator can establish Service Level Agreements (SLAs) and monitor quality of services relative to objectives found in the SLAs. Whenever a component of the infrastructure becomes unavailable, the IT administrator has immediate insight into which services and customers are affected.

When Insight Manager 7 and OpenView are linked together, they provide integrated management from the system to the service level. For example, the customer can go from managing their entire IT infrastructure from a service perspective with OpenView, to drilling down to investigate a component issue on an individual server with Insight Manager 7 and the Management Agents. If the server contains Lights-Out technology, the administrator has the control to resolve the issue remotely and immediately using the Lights-Out technology, Insight Manager 7, and other ProLiant Essentials software.

Dynamic resource optimization

Dynamic resource optimization, the third core capability, is the real-time delivery of computing resources whenever and wherever needed. The ability to match computing and storage resources to service demands reduces system cost by consolidating resources and simplifying management of complex environments. It improves operational efficiency by providing better overall resource utilization, response time, and capacity where needed.

HP technologies that enable resources to be optimized dynamically include:

- Scale-up and scale-out architectures,¹⁰ including modular blade architectures
- Workload optimization
- Resource deployment and provisioning

Scale-up architectures

HP's extensive product offerings allow customers to choose the appropriate application, OS, and platform to meet any service requirement. HP offers the most comprehensive scale-up roadmap in the industry. No other hardware vendor offers customers the flexibility to scale up from dense 2P rack-mount servers all the way to 64-socket modular architectures using industry standard IA-32 and Intel Itanium Processor Family processors.

Building on the success of the Profusion 8-way architecture, HP has designed its second generation of 8 way servers with the F8 chipset. The evolutionary F8 architecture obtains even higher performance than the Profusion architecture by optimizing the crossbar switch component and increasing bandwidths to match the processing power of the Intel Xeon MP processor. The high-performance 8-way architecture reduces potential performance bottlenecks through its innovative crossbar implementation and increases system availability with Hot Plug RAID Memory. The advanced F8 chipset allows HP to provide an industry-standard 8-way server that simplifies IT administration and lowers the total cost of ownership.

In 2003, servers based on the next-generation Itanium processor will scale up to 8-, 16-, 32 , and 64-way using a cell-based architecture (each cell contains four processors and associated memory modules). The Itanium-based systems will support three major operating environments: Microsoft® Windows, Linux, and HP-UX. These systems will deliver unprecedented performance within industry-standard architectures. In addition, these servers will provide a hardware partitioning ability previously unavailable in an industry-standard server configuration. The ability to hardware partition the server will allow customers to allocate resources as needed and to consolidate older, legacy systems into a single system at a single location.

¹⁰ Scale up refers to expanding an individual system's computing power by increasing processing power, memory, bandwidth, etc.

Scale out refers to expanding a solution's capacity horizontally by incrementally adding more servers, complete with processors, storage, and bandwidth, and sharing the user load across these servers.

Scale-out architectures

HP has leading scale-out architectures that provide the performance needed to meet demands ranging from web farms to technical computing. These architectures include:

- Scale-out architectures such as DISA
- Modular ProLiant BL blade architectures high-performance clusters
- Shared storage
- Fabric-computing solutions

Multi-tier, scale-out architectures such as DISA provide for resource flexibility by combining the best of both scale-up and scale-out into a comprehensive architecture for application services. The DISA structure results in an extremely stable IT environment with reduced risk of unplanned downtime.

Because it is a distributed architecture based on multiple tiers, the IT administrator can add compute, storage, or network resources incrementally and dynamically.

HP optimized the modular ProLiant BL server blade family for customers operating in a multi-tiered IT environment such as DISA. The ProLiant BL e-Class System contains a single-processor server blade that is ideal for specialized infrastructure servers such as firewalls or load balancers, for individual Apache or IIS web servers, or as a computational cluster node for the access layer of a DISA structure. It provides an ultra-dense design with up to 20 server blades per server blade enclosure. The ProLiant BL p-Class System, with its dual-processor, high-availability BL20p server blade, is ideal for use in mid-tier server applications such as dynamic web hosting, terminal server farms, or media streaming. The BL p-Class System has been developed with a flexible, intelligent, infrastructure that accommodates the existing server blades as well as future generations of two- and four-way server blades. The next-generation, two-way and four-way blades include fibre-channel capabilities to enable Storage Area Network (SAN) connectivity. The blades support HP and some third-party SAN storage. The ability to connect to SAN storage allows traditional enterprise applications to be deployed in a blade-only environment.

A customer desiring to move into a high-performance technical computing cluster can use ProLiant systems to scale up to hundreds of nodes. HP provides high-performance technical computing cluster solutions that run on ProLiant DL320, DL360, DL380, or BL e-class servers, using the Linux OS and third-party clustering software such as the Scyld Beowulf cluster¹¹ solution. Compute nodes are connected with high-speed interconnects and are devoted to compute-intensive parallel applications such as finite element modeling, geological mapping, and financial analysis.

For shared storage among servers and clusters, the HP family of Smart Array products uses common data formats, management tools, and components, greatly simplifying ongoing storage management and maintenance. This enables storage to be deployed as dedicated, clustered, or shared, and allows customers to migrate easily from direct-attached storage to storage area networks as their needs for capacity, performance, or availability change. For instance, an administrator might initially deploy an application on internal storage using a Smart Array Integrated 5i controller. If increased performance and capacity were needed, the administrator could migrate the drives to a Smart Array 5304 controller and Storage Works 4314 enclosure. Later, if even greater system availability were desired, the same drives could be moved to a Smart Array Cluster Storage device, using either Recovery Server Option or a ProLiant Cluster to improve system availability. All this can be accomplished without a single backup or restore process, since the drives can be migrated, even out of order, to each successive storage system. This versatility allows businesses of all sizes to enjoy the benefits of shared storage with the affordability and data protection capabilities of Smart Array controllers.

¹¹ For more information about high-performance clusters, see

<http://h18000.www1.hp.com/solutions/enterprise/highavailability/linux/hp-description.html> and <http://www.scyld.com/>.

Finally, an area that will be critical to the broad adoption of future scale-out architectures is the emergence of fabric-computing solutions, such as RDMA over TCP/IP protocols.¹² These RDMA technologies leverage existing IP networks, improve infrastructure utilization, and increase overall solution performance. As the datacenter becomes increasingly complex, the need for fabric infrastructures with pay-as-you-grow performance and a low-latency, high-bandwidth interconnect will be increasingly important. HP is leading the development of standards-based fabrics to enable these upcoming scale-out advancements.

Workload optimization

With the relatively low cost of industry-standard hardware, administrators are prone to simply add another server whenever application demands exceed the ability of a particular server. As a result, companies have to manage increasingly large server networks and deal with the accompanying issues of increased complexity and management overhead. By consolidating applications from two or more underutilized servers onto a single server, an IT administrator reduces cost and complexity, and uses the computing resources more efficiently.

The ProLiant Essentials Workload Management Pack (WMP) provides the tools a customer needs to fully utilize a server's processor and memory resources. With the WMP software, Resource Partitioning Manager (RPM), an IT administrator can configure resource partitions, that is, allocate processor and memory resources to specific applications. RPM reallocates resources dynamically to improve utilization and provide the stability to run multiple applications on a single server. It is a specific solution, based on Microsoft job objects, that allows workload consolidation on ProLiant servers running Windows 2000 family and Windows 2003 family operating systems. Workload consolidation reduces the total cost of ownership for servers and IT infrastructure by reducing IT administration costs, reducing software licensing fees, and reducing overall hardware costs.

HP ProLiant servers fully support Microsoft Windows Resource Manager (WRM) as another workload consolidation option for customers that are using Microsoft .NET. This job objects software is similar to RPM in that it pools resources and allows applications to have certain amounts of processors or memory allocated to them.

While some enterprises may wish to consolidate multiple applications (running on the same OS) onto a single server, others may want to consolidate multiple operating systems onto a single server. HP is partnering with VMware, Inc. to provide such a solution. VMware virtualization software¹³ consolidates and partitions servers for use with multiple operating systems. Applications and infrastructure services running on dedicated servers can be consolidated onto fewer, highly scalable, highly reliable enterprise-class servers. HP ensures that VMware software integrates seamlessly with Insight Manager 7 and other management software. By providing a seamless interface to ProLiant and industry-standard management tools, HP ensures that a server consolidation project will bring a high return on investment, saving power, real estate, and cooling costs.

Virtualization

In an adaptive infrastructure, an administrator should be able to view the infrastructure as a utility resource that provides compute, storage, network, software, and other capabilities. Resources should be dynamically and automatically assignable, with the ability to allocate resources easily and with minimal impact on applications. Virtualization transforms the physical characteristics of resources, providing a consistent logical view of the environment that enables simplified management and flexible asset provisioning. Adding a layer of software abstraction to resources enables IT administrators to more easily match enterprise resources to service demands in real time.

ENSAextended is the HP storage strategy that provides a blueprint for creating a virtualized storage infrastructure. Storage management applications, such as HP OpenView Storage Area Manager,

¹² RDMA is an acronym for Remote Direct Memory Access. TCP/IP is an acronym for Transmission Control Protocol/ Internet Protocol, the most common transport layer protocol, used on Ethernet and the internet.

¹³ For more information about VMware, visit their website at <http://www.vmware.com/>.

view the storage as a utility, or virtual pool, while virtualization software manages the underlying physical devices. This virtualized storage architecture can be one component of a more fully virtualized infrastructure such as the HP Utility Data Center.

The HP Utility Data Center with Utility Controller Software provides a solution for planning, designing, managing, provisioning, and supporting data center resources. Using the Utility Controller Software, administrators can create and run IT environments as a highly automated service. It allows an enterprise to establish the original infrastructure, then reallocate server, network, storage, and related resources through the advanced data center management software. As a result, IT administrators can respond quickly as business needs change.

Resource deployment and provisioning

Computing resources are often labor intensive to configure and deploy. Once a resource is deployed, the IT administrator must face another labor-intensive process to maintain the OS, device drivers, or application software when updates to these become available. If the administrator delays an update or does not apply it in a systematic manner, both security and system availability may be affected. As the demands of the business or of particular applications grow, the administrator must determine how best to meet those changing demands for compute, network, or storage resources.

HP deployment tools allow a customer to deploy a single server or thousands of servers in a rapid manner. The SmartStart CD provides a software tool for simplified and reliable single-server deployment. The SmartStart Scripting Toolkit¹⁴ simplifies multi-server and Smart Array storage deployment by providing administrators a set of utilities for automating the configuration process. By incorporating scripts for hardware configuration into the SmartStart Scripting Toolkit process, IT administrators can rapidly and remotely configure a new server, install the OS and configure the Smart Array storage system, making it possible to scale server deployments to large volumes with minimal effort.

The ProLiant Essentials Rapid Deployment Pack builds upon the technology of the SmartStart Scripting Toolkit to scale server deployment to high volumes. It maximizes a customer's resources by providing a full server build from a remote console, automated server configuration on the fly, and installation of standard software sets based on server functions.

Furthermore, the Rapid Deployment Pack is optimized to enable "rip and replace" functionality for the modular ProLiant BL server blades. An administrator can use the tools to pre-assign a particular function to each server blade slot in an enclosure. If a particular server blade were to fail at some later time and be replaced, the new server blade would automatically take on the role of the failed server blade, significantly reducing the time and effort needed to keep servers in production. For a service provider environment, quick deployment and redeployment are essential. The ProLiant BL servers contain embedded firmware to support remote ROM flash, which simplifies system maintenance. The ROM flash capability is scriptable for unattended operation and compatible with PXE-based¹⁵ deployment and management tools. Using the previously mentioned automated software and hardware deployment tools, ProLiant BL server blades can be rapidly provisioned or redeployed, including the hardware configuration, OS, and applications.

Planning for upgrades or changes to the IT infrastructure is an essential component of resource deployment and provisioning functions. An administrator must accurately assess the hardware and software needed to meet critical business demands. To assist with this process, HP has planning tools¹⁶ such as System Configurator to help create configurations and bills of material for ProLiant servers and storage. HP also has extensive, web-based application sizers to help the customer

¹⁴ The Smart Start Scripting Toolkit supports the ProLiant DL and ML families of servers, as well as some earlier ProLiant servers. More information can be found at <http://h18000.www1.hp.com/manage/deployment.html>.

¹⁵ PXE is an acronym for pre-boot execution environment. PXE enables servers to load and execute a network bootstrap program and a pre-configured image of an OS or boot floppy.

¹⁶ For more information about HP planning tools, see <http://activeanswers.compaq.com>.

determine the size and scope of the solution environment and a full complement of solution guides specific to business needs. Use of these tools allows a customer to deploy and operate an IT environment based on best practices, thus ensuring quality of service, fast solution time, and maximum operating efficiency.

Open industry standards

The core technologies of the adaptive infrastructure are based on cost-efficient, open industry standards. Customers can be confident that an infrastructure based on unifying industry standards will protect their capital investment in technology, will have the full support of independent software and hardware vendors, and will have best-in-class industry expertise available. Industry-standard architectures are becoming pervasive in three key functional areas:

- Hardware architectures
- Fabric interconnects
- Software interfaces

By developing innovative hardware platforms, interconnects, technologies, and software that are based on the framework of industry standards, HP is able to focus on delivering an IT infrastructure that is standard, pervasive, cost efficient, and flexible.

Hardware architectures

HP is building the next generation of industry-standard servers on IA-32 and IPF processors. The IA-32 processor family will continue to provide a cost-effective, stable, powerful architecture for the majority of industry-standard servers. HP is consolidating its entire 64-bit server family on the Intel Itanium architecture to provide customers unparalleled price:performance in high-performance servers. HP is the only server provider that can service any IT environment with a complete line of IA-32 and IPF-based servers. Because HP has focused its server design efforts on a single architectural family, HP will propel innovation within server design and promote innovation among other vendors.

Fabric interconnects

Fabric interconnects open the possibility of unified networking, storage, and cluster communications over a common type of interconnect. Use of a converged fabric infrastructure has the potential to greatly simplify the datacenter infrastructure. HP is focused on developing fabric technologies that extend the capabilities of existing datacenter networks and provide straightforward infrastructure transitions. These fabric standards will provide cost-effective, low-latency fabric interconnects to enable scalability throughout the datacenter.

For example, the Internet Small Computer Systems Interface (iSCSI) protocol¹⁷ provides block-level data transfer over ubiquitous Ethernet networks, which use the IP protocol. By leveraging Ethernet and IP standards, iSCSI can potentially extend storage accessibility from SANs to local area networks, wide area networks, or to a global storage network. Other advances in Ethernet networking technologies, such as the recently completed RDMA over TCP/IP standard, will drive extremely scalable performance by improving communication efficiency and reducing latencies. RDMA over TCP/IP allows servers to communicate without interim data copying operations (zero copy), thus reducing memory bus bandwidth demands and improving overall server performance. RDMA/TCP improves solution availability, as memory and processor resources are free to do more useful work.

¹⁷ For more information about iSCSI, see "iSCSI Technology: A Convergence of Networking and Storage," available at <http://h200001.www2.hp.com/bc/docs/support/SupportManual/c00256917/c00256917.pdf>.

Software interfaces

The use of standard software interfaces, for example, Extensible Markup Language (XML) and Simple Object Access Protocol (SOAP), provides a common framework for intelligent middleware applications such as management tools routinely used in datacenter environments. Standard software interfaces will enable multiple tools within a datacenter to share information seamlessly, to be accessible through standard web browsers, and to be more efficient at integrating complex processes.

Integrated with solutions and services

A key attribute in HP's delivery of an adaptive infrastructure is the ability to integrate into the IT infrastructure solutions and services that bring lasting value. HP recognizes that customers are looking for a step-wise approach to achieving greater agility. HP's portfolio of infrastructure solutions combines products, services, and partnerships to enable customers to meet their critical needs today while deploying increasingly flexible and efficient infrastructures.

Adaptive infrastructure services

The HP services global team collaborates with customers to enable services that reduce the risk and complexity of managing large datacenter environments and increase business agility. Furthermore, HP services works carefully to align closely with the customer's evolving business needs and to bridge the gap between business needs and IT capabilities.

HP services helps deliver the adaptive infrastructure through the following services:

- Adaptive infrastructure services — strategy, assessment, and architecture services to determine optimum solutions and prioritize improvement opportunities for a customer's specific infrastructure needs
- Consulting and integration — services to architect, test, and implement recommended adaptive infrastructure practices
- Customer support — support services integrated with automated provisioning and management to enable the complete implementation of a customer's adaptive infrastructure
- Managed services — dynamic outsourcing to improve business agility, extend expertise, and extend financial flexibility for customers

The goal of the HP services global team is to build an IT environment that meets the customer's current requirements, protects existing investments, and yet allows room for future innovative technologies.

Adaptive infrastructure solutions

HP provides infrastructure solutions that include:

- IT consolidation solutions — optimizing processes and system to be more efficient and effective for customers who wish to restructure their existing environments to conserve costs, yet also require increased adaptability for the future.
- On-demand solutions — integrating products, services, and financing to deliver IT resources when and where a customer needs them. These solutions allow for pay-per-use, pay-per-forecast, and managed alternatives for customers who want to match their investments and cost profiles to the use of their technology.
- Business continuity solutions — maintaining uptime in mission-critical environments and providing the ability to recover quickly from disruptions and disasters.

- Management and security solutions — reducing complexity and improving fault-tolerance in the customer's current environment by incorporating technologies such as Insight Manager 7 and OpenView.
- Business integration solutions — enabling customers to maximize the interoperability of their business solutions and delivering the technologies needed to implement robust, integrated environments such as business-to-business and business-to-customer communities.

Collaborative industry partnerships

HP's partnering strategy — developing cooperative relationships with key solution, software, and hardware providers — is fundamental to delivering the highest value and choice for customers. By taking a collaborative approach to development, HP is able to focus its research investments in areas that do not compete with partners that are industry leaders in other areas.

HP works with industry leaders like Microsoft, Intel, Red Hat Linux, Novell, and others to develop completely integrated, scalable solutions based on ProLiant technologies. Because ProLiant servers are the platform of choice for a variety of business applications such as SQL Server, Oracle, and Siebel, ProLiant servers and storage solutions often become the development platform for new generations of these applications. HP engineers work closely with third-party software and hardware providers to test and optimize the solutions. This collaborative approach results in more robust and better integrated solutions for customers.

As an example, HP has an extensive Developer and Solution Partner Program (DSPP) with over 6,000 members, including systems integrators, independent software vendors, independent hardware vendors, and developers. In this program, HP provides business development resources, developer resources, specialized programs to help build customized solutions, and training to developers and partners. Within DSPP, the HP blade server alliance program has been established to assist partners with developing and marketing their software solutions on HP ProLiant BL blade servers. The result is software and solutions optimized for the blade server environment and an accelerated time to market for these optimized solutions.

Conclusion

HP's vision for an adaptive infrastructure is based upon three industry-defining core capabilities: continuous, secure operations; automated, intelligent management; and dynamic resource optimization. In each of these core technology areas, HP has established hardware and software solutions that enable an adaptive infrastructure.

The HP adaptive infrastructure is not just a goal for the future. HP server, storage, and software products already provide functionality that a customer can implement incrementally along a logical growth path. ProLiant technologies and solutions are aligned to the adaptive infrastructure vision and provide an excellent standards-based foundation for the adaptive infrastructure. The net effect of the adaptive infrastructure will be an IT organization that can adapt quickly to changing business conditions, conserve valuable IT resources, and provide the highest level of customer service.

Call to action

To help us better understand and meet your needs for ISS technology information, please send comments about this paper to: TechCom@HP.com.

© 2003 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

HP-UX Release 10.20 and later and HP-UX Release 11.00 and later (in both 32 and 64-bit configurations) on all HP 9000 computers are Open Group UNIX 95 branded products.

Intel Itanium Processor Family is a trademark of Intel Corporation in the U.S. and other countries and is used under license.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corp.

Oracle is a registered U.S. trademark of Oracle Corporation, Redwood City, California.

TC030202TB, 02/2003

Printed in the US

